

Defending Critical Infrastructure

Gerald Brown, Ph.D.
Matthew Carlyle, Ph.D.
Javier Salmeron, Ph.D.
Kevin Wood, Ph.D.

Operations Research Department
Naval Postgraduate School
Monterey, California 93943
<http://www.nps.navy.mil/or/roster.htm>

17 December 2004

The attacks of September 11, 2001 and subsequent events have motivated the United States to study its critical infrastructure anew, and to seek improvements that make it more resilient to terrorist actions. As operations researchers, we develop mathematical models and analytical tools to study organizations and guide their infrastructure and operational improvements. This white paper shows how we have brought the analytical techniques of OR to bear on the specific problem of protecting critical infrastructure from attack. Our mission is to strengthen societal continuity through dangerous times by 1) making critical infrastructure more resilient to attack, 2) helping governments and corporations plan the improvements that will provide that resilience, and 3) influencing public policy regarding, for example, investment incentives for hardening critical infrastructure. Our government has formed the Department of Homeland Security (DHS), and our Department of Defense now has the new Northern Command (NORTHCOM 2004); both are tasked with defending our country inside US borders. This is an important problem.

Analysis

What is critical infrastructure? The National Strategy for Homeland Security (DHS 2002) deems 14 systems critical to the United States (Table 1). These include entities such as “Government” and “Public Health,” but most systems on this list have, at their foundation, physical infrastructure that connects components of our economy and transports among these components the materials and information that define our economy. These infrastructures include, but are not limited to, road, air, rail, telecommunications, electric power, fuel, and water.

Agriculture
Food
Water
Public Health
Emergency Services
Government
Defense Industry
Information and Telecommunications
Energy
Transportation
Banking and Finance
Chemical Industry
Postal and Shipping
Key Assets

Table 1. **Fourteen infrastructures critical to the United States, as defined by the Department of Homeland Security (2002).**

These infrastructures represent huge investments of our nation’s wealth. Minor disruptions to components in any one infrastructure can have a severe impact on its performance, such as a massive power outage resulting from losing a few key subsystems. Disruption of one system can also interfere with operation of other infrastructures, such as when electric power outage impacts telecommunications.

Most of these infrastructures have been engineered and built to withstand potential disruptions due to accidents or random acts of nature. A typical vulnerability assessment for such a system considers the removal of one component at a time or one act of nature at a time, and pronounces the system robust if there is no “single point of failure” that is crippling. Frequently, the end result of designing a system based on such analysis is one with highly-reliable components protected by cheap shutoff systems. The US electric power transmission systems are the best-known examples of this.

It is no longer sufficient to merely harden these infrastructures to resist single points of random failure. An intelligent attacker can see what is vulnerable as well as we can. For example, a lone attacker with a high-powered rifle can cause grave damage to an entire electric power grid by targeting highly reliable components --- components for which spares are rarely needed --- at a few remote, vulnerable, undefended substations.

A captured Al Qaeda training manual (Department of Justice, 2004) advises: “Using public sources openly and without resorting to illegal means, it is possible to gather at least 80% of information about the enemy.” Our experience is that this is a conservative estimate. A malicious, intelligent adversary has the ability to observe the weaknesses in our infrastructures, and exploit them.

To see how an intelligent enemy might evaluate our infrastructure, we draw from US military doctrine on defending such assets. Table 2 shows the doctrinal components used by the US Army to determine the value of defended assets.

Criticality	How essential is the asset?
Vulnerability	How susceptible is the asset to surveillance or attack?
Reconstitutability	How hard will it be to recover from inflicted damage, considering time, special repair equipment, and manpower required to restore normal operation?
Threat	How probable is an attack on this asset?

Table 2. **Criteria for assessing the value of defended assets (Department of the Army 2002a,b).**

Intelligence agencies provide many of these assessments, which are, of necessity, highly subjective. For instance, how “reconstitutable” is the Lincoln Memorial? What is the threat it will be attacked? Of course, the threat, or probability of an attack, is decided by our opponent: We must plan for what is possible, rather than what is likely.

We can add objectivity to these assessments. We roll up all constituent considerations into one gauge: **criticality**. In our lexicon, “criticality” may include consideration of vulnerability, reconstitutability, and threat, but we use criticality as a quantitative indication of how the loss of a component or components affects the performance of the infrastructure.

Al Qaeda teaches as its primary mission “overthrow of godless regimes (by) gathering information about the enemy, the land, the installations, and the neighbors, ... blasting and destroying the places of amusement, ... embassies, ... vital economic centers, ...bridges leading into and out of cities, ...” (Department of Justice, 2004). We distill from all of this that their target value assessments are exactly the same as ours. After all, their objective is to maximize damage to us as we feel it. They want to maximize whatever it is that we want to minimize. Accordingly, we assume that our target value assessments are shared by our enemies.

These value assessments must also be made for any *set* of defended components for which a joint attack would inflict super-additive damage, or collateral damage to systems not attacked at all. Hereafter, we assume that engineers and other analysts have made acceptable assessments of the values of individual infrastructure components, and have determined how those components interact in the infrastructure as a whole. This will allow us to assess the value of *sets of components* subject to simultaneous attack and damage. We also assume that we can estimate the “cost” of any candidate attack to the attacker, and any other reasonable restrictions on his ability to mount attacks.

What We Have Done

In response to the insight that a belligerent enemy can learn essentially everything we know about our own critical infrastructure and its criticality, we have developed decision support tools that explicitly mimic an informed attacker: someone who knows the structure of the system he is attacking, and who can predict how the owner of that system will respond to attacks. With such a tool, we can predict just how an attacker can determine the most damaging attacks (possibly coordinating several simultaneous attacks), and can prescribe the resulting reaction of the system and its operators. (See Figure 1.)

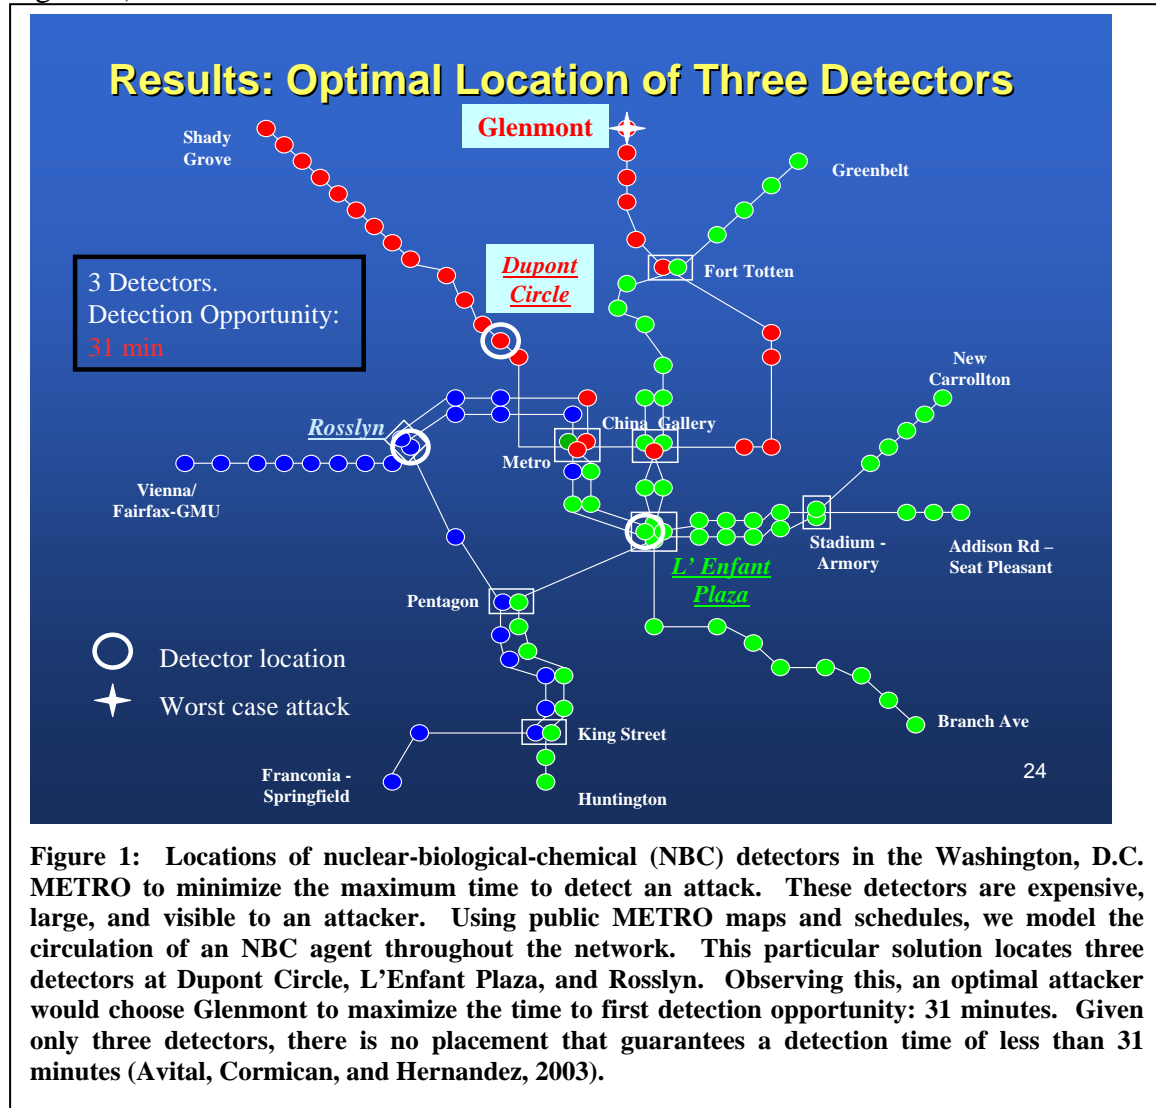
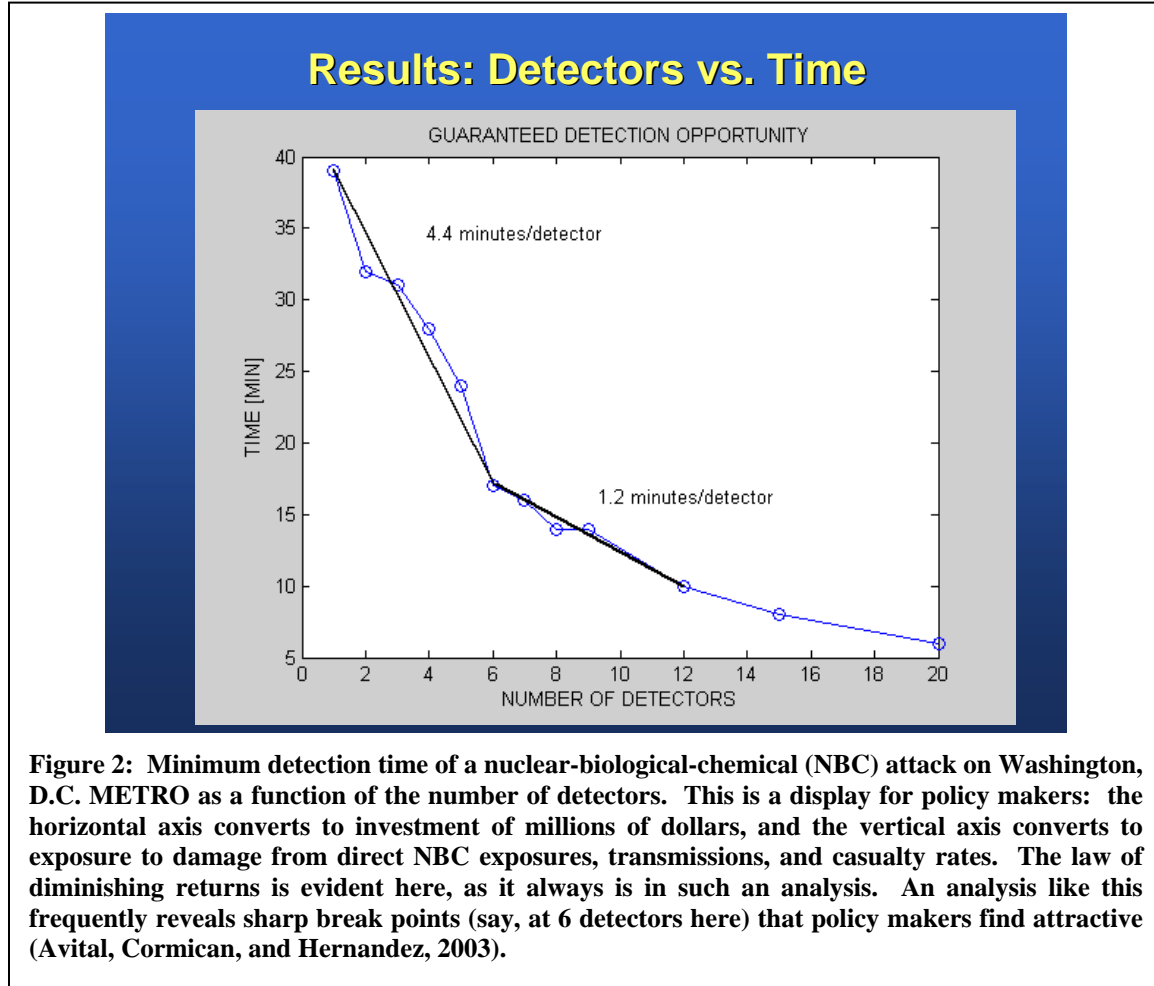


Figure 1: Locations of nuclear-biological-chemical (NBC) detectors in the Washington, D.C. METRO to minimize the maximum time to detect an attack. These detectors are expensive, large, and visible to an attacker. Using public METRO maps and schedules, we model the circulation of an NBC agent throughout the network. This particular solution locates three detectors at Dupont Circle, L'Enfant Plaza, and Rosslyn. Observing this, an optimal attacker would choose Glenmont to maximize the time to first detection opportunity: 31 minutes. Given only three detectors, there is no placement that guarantees a detection time of less than 31 minutes (Avital, Cormican, and Hernandez, 2003).

These decision support tools rely on models based on our experience as military planners, targeting enemy infrastructures. When we turn the analysis back on our own systems we discover how vulnerable they are to attack, and we can investigate how to harden our systems to reduce the damage of a “worst-case” attack by an intelligent adversary through such actions as hardening single components, adding the right level of

redundancy, defending key assets, detecting an impending attack, etc. (See Figure 2.)



We have built models of most of the infrastructures mentioned above. To test each of these, we assemble a “red team” of well-trained military officers to gather each case scenario from strictly public sources --- we use no government identification or privileged access of any kind, and take care to leave no trace of our investigation. We use each hypothetical scenario to plan both the role of the attacker and of the defender.

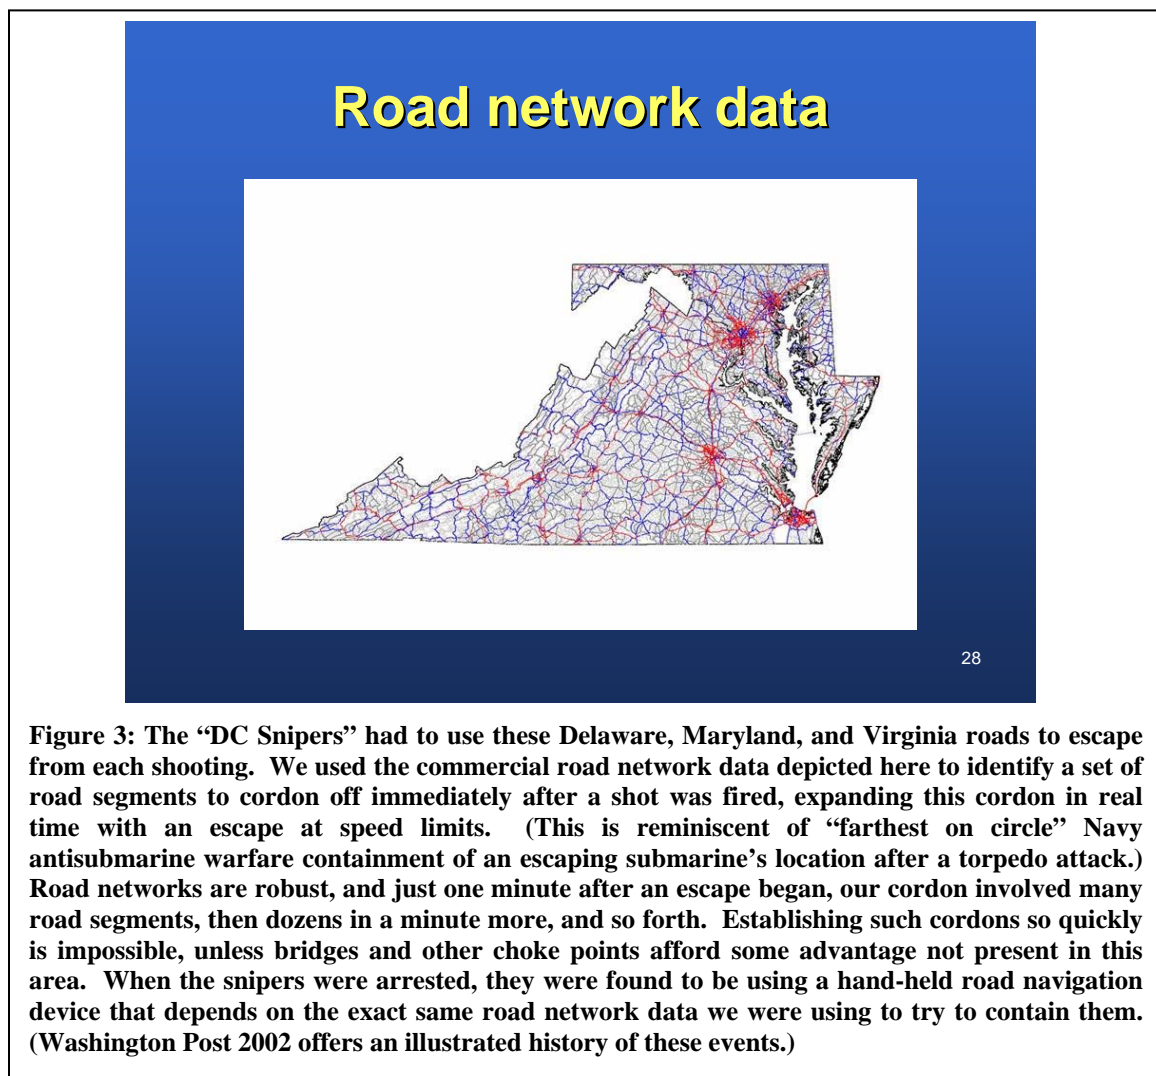
In almost every case we find that the attackers can do more damage than we would have predicted, for various levels of effort, and that their optimal attacks frequently do not target the “obvious” components revealed as critical in “single point of failure” analyses. We also discover valuable insights about defending or hardening our systems after we see what optimal attacks look like, for varying levels of effort by attacker and defender.

We have found that the attacker, even if he broadcasts his intentions to his victim, usually possesses a tremendous advantage. This is the reverse of classical military theory, and accrues from the hugely asymmetric nature of this conflict: the defender must protect a huge, dispersed target set, while the attacker need only focus on his small

set of targets chosen to maximize damage effects.

We have also found time after time that building robust systems, or hardening the ones you already own, can be very expensive. However, if you understand what the optimal attacks must look like, you can make your system much more robust, for the same investment, than if you harden it based on “single point of failure” analysis.

It turns out that our road systems are remarkably robust (see Figure 3), that fuel distribution systems are unbelievably fragile (see Figure 4), and that most other systems lie somewhere in between. We have also discovered that every infrastructure has a breaking point, beyond which damage is catastrophic. Knowing that, we try to learn what we might do to maximally delay such a breakdown given some affordable budget.



ELECTRIC POWER GRID VULNERABILITY

Javier Salmeron and Kevin Wood

Sponsored by Department of Justice

Grid Snapshot

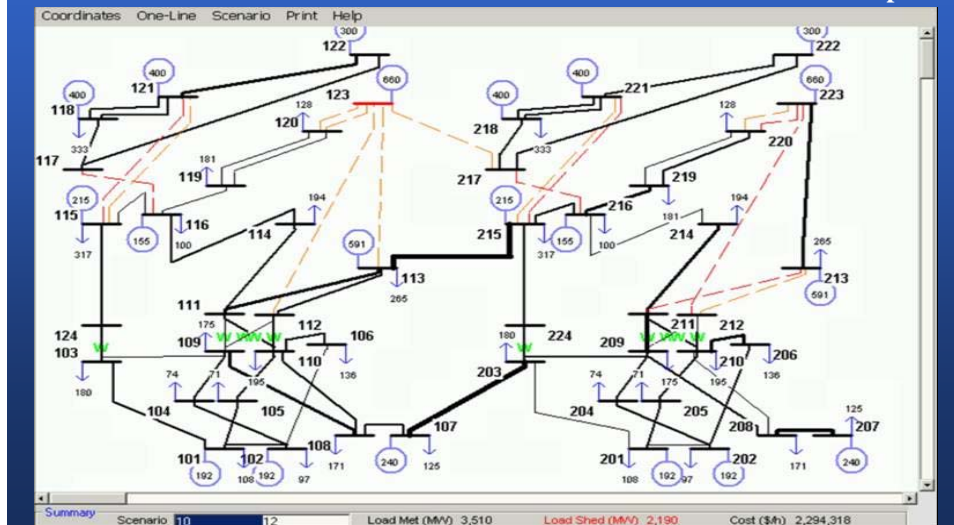
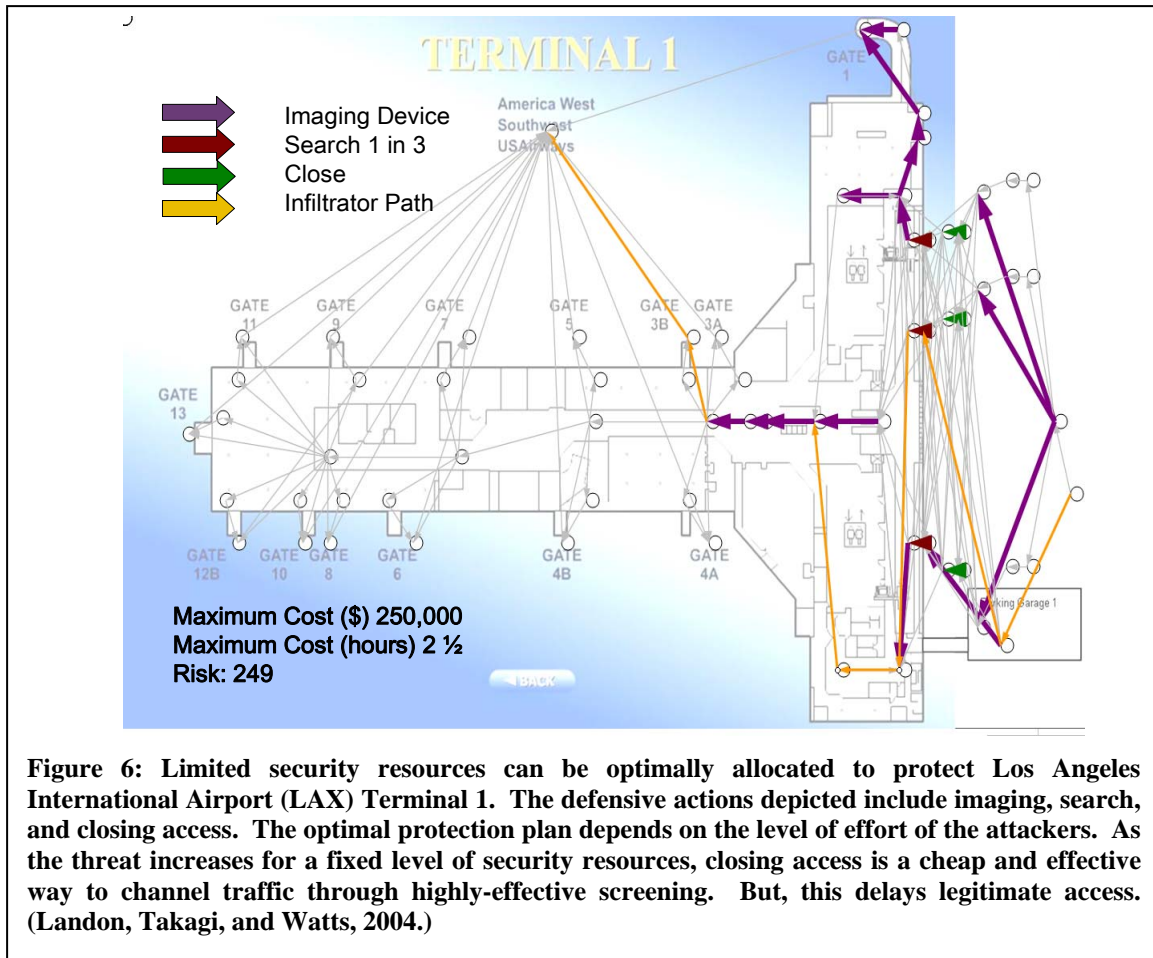


Figure 5: Vulnerable Electric Grid Analyzer (VEGA) screen shot. The icons depict generators, transformers, transfer bars, transmission lines, and customer demands. For any given level of attacker capability, VEGA finds an optimal target set of critical vulnerable components. The dashed lines show where capacity has been lost due to such an attack, despite optimal response by the system operator to bring reserve generation capacity online, re-route power and reduce deliveries of interruptible customer demands. As attacker capability increases, VEGA discovers a succession of most-critical component sets to attack, and shows how the entire power grid can best respond to this damage. An animation of this leads to non-intuitive insights: this is far more subtle than simple “single point of failure analysis,” and the results aren’t obvious (Salmerón, Wood, and Baldick, 2003).

Protection of critical assets requires planning how to detect and/or deter attackers or insurgents. Although this has long been a topic of interest to our military, we have not traditionally granted our enemies the advantage of seeing our defensive preparations. We have had to revisit this topic to see how our plans are influenced when we know our enemies will observe our preparations (see Figure 6.).



Supply chains --- physical distributions systems --- are the key infrastructure of most private-sector companies that manufacture or distribute goods. Strategic supply chain design has a long and successful record in the US, reducing costs and increasing service levels. Unfortunately, efficient supply chains are very fragile. In fact, after you scrupulously invest exactly the right amounts of money on exactly the right bottlenecks in your supply chain, you create product flow patterns that resemble spanning trees. A spanning tree is a maximally-fragile network: Break any link, and the network is disconnected.

We have teamed with Prof. Terry Harrison of Pennsylvania State University and Dr. Jeffrey Karrenbauer, President of INSIGHT, Inc., a company devoted to supply chain optimization for more than 25 years (INSIGHT 2004). With them we have analyzed detailed corporate supply chain data for scores of companies, including the majority of the FORTUNE 50. We have collaborated in designing new features for INSIGHT's supply chain optimization tools to evaluate supply chain vulnerability, and to optimally determine what to do to harden these supply chains (many key results have already been presented by Brown et al. 2003a,b, Brown, Carlyle, Harrison, et al., 2004).

Surprisingly, you can achieve a great deal with a modest investment in planning

and additional physical infrastructure. Sometimes, just retaining and strategically relocating some spare capacity that already exists can have a beneficial effect at virtually no incremental cost.

We are presenting our findings to any company that invites us, and we are gratified by the enthusiastic response we get to relatively simple discoveries. American companies now have senior executives with titles and position descriptions focused on “preserving corporate continuity.” These were originally motivated by threats to ubiquitous information systems: Back-up computer facilities and critical data storage have been in vogue for some time. Now, these same companies are coming to recognize that they also need to attend to their ability to continue physical operations after denial of access to some infrastructure they depend on, whether or not the damage was aimed at them, or they own the infrastructure they depend on.

Paradoxically, fewer than half of US companies have any form of terrorism insurance (Fleming 2004). Yet, since 9/11 the magnitude of recognized exposure has increased, and one would think companies should address this new level of risk. The Terrorism Risk Insurance Act (United States Public Law 2002) is a stopgap law to help companies redesign their insurance plans for our newly-dangerous world: This law expires in 2005, and the future of this legislation is uncertain. Of course, financial losses are merely a telltale of deeper damage to a company whose continuity of operation is interrupted.

We still encounter considerable confusion in the private sector between random acts of nature --- these have been studied by insurance actuaries for centuries --- and belligerent acts of intelligent terrorists who observe defensive preparations and intend to maximize damage. Our case studies illustrate the qualitative distinction between these types of threats.

We have also learned from these companies that aggrieved labor unions and rapacious competitors can be just as clever and determined as terrorists, and have exactly the same goals: to maximize damage (to market share, or to profit, etc.) inflicted. The denial of access to our West coast ports in 2002 due to a labor dispute was no less damaging than the anthrax attacks of 2001 that closed postal and shipping services on our East coast. When you model your supply chain and pose such modal threats, it becomes clear how and where to invest in a little more capital, or to limit exposure to such attacks, to achieve capacity that resists such attacks.

This work has also yielded some new military and diplomatic planning tools. We have developed a decision support tool to plan theater ballistic missile defense (Brown Carlyle, Diehl, Kline and Wood, 2004), where the attacker can see some or all of our defensive preparations, and another to discover how best to delay a covert nuclear weapons program (Brown, Carlyle, Harney, Skroch, and Wood, 2004), where the proliferator will observe some or all of our actions to keep him from completing a finished weapon.

A key insight from these military and diplomatic exercises is that deception and secrecy can make huge contributions to defending critical infrastructure, but that we cannot estimate the value of secrecy and deception until we understand the worst case attack in the completely transparent case.

Although this work is all relatively new, there is already an emerging body of unclassified publications including about fifty case studies, several graduate theses, open-literature publications, and a number of prototypic decision support tools. Table 3 shows some of these case studies. We are working with the institutions that plan for dealing with these threats and welcome inquiries. We also provide classified products to planners as the need arises.

Electric grids Road networks Strategic rail network Domestic water system in Southern California Sea lanes, canals, and restricted straits Multicommodity supply chains Petroleum distribution network in US Southwest; Northern California; and Defense Fuel Supply System, Japan Weapon of mass destruction (WMD) at Super Bowl, Reliant Stadium, Houston, Texas; Washington, DC, metro; Meeting of heads of state, Melbourne, Australia; Changi Naval Base, Singapore; Manhattan; and Norfolk, Virginia Insurgent incursions DC sniper escapes Leontief economic attack WMD development project Theater ballistic missile attacks
--

Table 3. Case studies such as these have evaluated roles of both attacker and defender.

What We Have Learned

The data is out there, and if we can get it, anybody can. “Sunshine laws” require that our governments, federal to local, conduct their affairs in as open and transparent a fashion as possible, and they do. The world-wide web makes any public posting a global one. Attractive web sites are universally fashionable, and government agencies and municipalities have produced lots of “cool web sites.” Many web sites have been redesigned in the last couple of years to make it harder to access key information, but we have found stunning exceptions. We have advised anybody who will listen to appoint an independent “red team” to analyze a public website with intent to plan harm to its sponsor.

The answers aren’t obvious. The most damaging coordinated attacks, or the

most effective defense, are often non-intuitive. It turns out that key US infrastructures are huge systems, and analysis at large scale deserves some fairly rigorous, purpose-built decision support tools to formalize the notion of a transparent, two-sided conflict. Manual analysis or simulation is better than no analysis, but in our work optimized solutions are the ones that bring the most surprises and lead to the deepest insights.

Fortifying infrastructure is expensive. Critical infrastructures have been built at enormous expense to be efficient. Efficient infrastructure is frequently fragile. Dealing with this may cost a lot of money. For those infrastructures owned and operated by private entities, there is no economic incentive to spend huge sums of money for this. This calls for government subsidies, changes to tax codes, and regulatory reform to create an environment motivating incremental fortifying investments.

Defending infrastructure is expensive. In military vernacular, the US is a target-rich country. Mounting a small-scale attack is cheap for an opponent. This is an asymmetric conflict.

Malicious, coordinated attacks are much more damaging than random acts of nature. This is the biggest and most dangerous misconception we encounter in our audiences. A small-scale attack can inflict more damage than a major hurricane, great earthquake, etc.

Reliability is not the answer. Single point of failure analysis is insightful, but the failures to worry about are those of the most critical, rather than the least reliable infrastructure. Many infrastructure designers confuse reliability with criticality, and this mistake completely conceals the fact that malicious, coordinated attacks on critical infrastructure may *target* the extremely reliable components for precisely the reason that this will inflict the most enduring damage.

The right redundancy may be the answer. For any given level of investment, there is usually a dominant set of incremental changes to infrastructure that return maximal immediate benefit. Often, a great deal of benefit can be achieved at relatively modest cost, by adding a few alternate shipment paths, or installing some excess capacity at just the right locations, etc.

Secrecy and deception can be valuable. There is good reason to keep your plans to yourself. It's not easy to keep major investments in infrastructure and defense secret. But, it's worth trying.

Finally, **worst case analysis using optimization** is key to a credible assessment of infrastructure vulnerability. We *cannot* perform single point of failure analyses and hope that we are adequately protected. We *cannot* assume attacks happen randomly. We are facing a determined, intelligent enemy who seeks to do us maximal harm. Knowing this, we can prepare for the worst case using tools from optimization, and, using those same tools, we can discover the most effective use of our defensive resources to thwart our enemy's plans.

Acknowledgements

Brown and Wood are grateful for sustaining research support by the Air Force Office of Scientific Research (Optimization and Discrete Mathematics Program), the Office of Naval Research (Division of Mathematical Sciences), and the Joint Warfare Analysis Center. Salmerón and Wood have been supported by the Department of Justice (Department of Homeland Security). Brown, Carlyle, and Wood are supported by the National Security Agency. At various times, we have also worked with and received support from every US uniformed military service. In late 2001, we approached INSIGHT, Inc., to help us discover what private companies could do to fortify their operations against hostile threats. (Note: Brown and Wood have worked on private-sector business optimization problems for decades with INSIGHT.) INSIGHT has granted unfettered use of its supply chain design software, devoted extensive development effort, provided data from a host of private-sector clients (scrubbed of proprietary confidential identification), and arranged direct access to its clients.

References

Andrews, LCDR Charles, USN, Capt Kemp Cason, USMC, MAJ Alison Godfrey, USA, and Capt Mark Revor, USMC, *Saudi Arabian Pipelines Red Team Report*, November 2003.

Avital, LCDR Ittai, Israeli Navy, LCDR Kelly Cormican, USN, LTC Andy Hernandez, USA, *Washington D.C. Metro Red Team Report*, November 2003.

Brown, G., M. Carlyle, D. Diehl, J. Kline, and K. Wood, 2004, "How to Optimize Theater Ballistic Missile Defense," *Operations Research*, (to appear).

Brown, G., M. Carlyle, R. Harney, E. Skroch, and K. Wood, 2004, "Interdicting a Nuclear Weapons Project." (in review).

Brown, G., M. Carlyle, T. Harrison, J. Salmeron, and K. Wood, 2004, "Designing Robust Supply Chains and Hardening the Ones You Have," *INFORMS Conference on OR/MS Practice*, Cambridge, MA, April 26 and 27.

Brown, G., M. Carlyle, T. Harrison, J. Salmeron, and K. Wood, 2003a, "How to Attack a Linear Program," plenary address, *Military Operations Research Society*, Quantico, VA, June 10.

Brown, G., M. Carlyle, T. Harrison, J. Salmeron, and K. Wood, 2003b, "Tutorial: How to Build a Robust Supply Chain or Harden the One You Have," *INFORMS Annual Meeting*, Atlanta, GA, October 19.

Department of the Army, 2000a, *Field Manual FM 3-01.11 Appendix A – ADA Employment Principles, Guidelines, and Priorities*, accessed 3 May 2004 at <http://www.globalsecurity.org/military/library/policy/army/fm/>

Department of the Army, 2000b, Army Field Manual FM 44-100 Chapter 4. Fundamentals of Army Air and Missile Defense Operations, accessed 3 May 2004 at <http://www.globalsecurity.org/military/library/policy/army/fm/>

Department of Homeland Security 2002, National Strategy for Homeland Security, <http://www.whitehouse.gov/homeland/book/> , accessed 2 December 2004.

Department of Justice, 2004, “Al Qaeda Training Manual”, <http://www.usdoj.gov/ag/trainingmanual.htm> , accessed 2 December 2004.

Fleming, C. 2004, “Terrorism Insurance: Many Companies Are Going Without,” **Wall Street Journal**, Dec. 13, p. C1-2.

Insight, 2004, Strategic Analysis of Integrated Logistics Systems (SAILS), Manassas, VA, www.insight-mss.com , accessed 2 December 2004.

Israeli, E. and K. Wood, 2002, “ Shortest-Path Network Interdiction ,” *Networks*, Vol. 40, pp. 97-111.

Landon, LT Chris, USNR Capt Koichi Takagi, USMC, CPT Krista Watts, USA, *LAX Terminal 1 Red Team Report*, November 2004.

NORTHCOM 2004, <http://www.northcom.mil/> , accessed 2 December 2004

Salmerón, J., K. Wood, and R. Baldick, 2003, “ Analysis of Electric Grid Security Under Terrorist Threat ,” *IEEE Transactions on Power Systems* (to appear).

United States Public Law 2002, “Terrorism Risk Insurance Act,” Public Law 107-297, 116 Stat. 2327, <http://www.gao.gov/new.items/d04307.pdf>, <http://www.treas.gov/offices/domestic-finance/financial-institution/terrorism-insurance/>

Washington Post 2002, <http://www.washingtonpost.com/wp-srv/metro/daily/oct02/snipershootings.htm>.

Wood, R.K., 1993, “ Deterministic Network Interdiction ,” *Mathematical and Computer Modelling*, 17, pp. 1-18.